

CLAIMS

1 1. An apparatus comprising:
2 at least one processor;
3 a memory coupled to the at least one processor;
4 a cluster engine residing in the memory and executed by the at least one processor,
5 the cluster engine providing a mechanism for communicating ordered messages to and
6 from a plurality of nodes in a computer cluster, wherein the apparatus comprises one node
7 in the computer cluster;
8 a protocol residing in the memory that specifies at least one data message and at
9 least one acknowledge (ACK) round that provides a time benchmark for determining
10 whether or not a message has been received without using any timer; and
11 a job residing in the memory and executed by the at least one processor, the job
12 processing the protocol, wherein the job functions according to receiver logic that uses
13 the at least one ACK round to determine without using any timer whether the at least one
14 data message has been received.

1 2. The apparatus of claim 1 wherein the protocol comprises a plurality of phases that
2 are each followed by an acknowledge (ACK) round.

1 3. The apparatus of claim 2 wherein each phase of the protocol is defined so that no
2 node can both send a data message and receive a data message during any phase of the
3 protocol.

1 4. The apparatus of claim 1 wherein the receiver logic comprises post-ACK logic
2 that determines whether the at least one data message has been received after the job
3 processes a selected one of the at least one ACK rounds.

1 5. The apparatus of claim 1 wherein the receiver logic comprises pre-ACK logic that
2 determines whether the at least one data message has been received before the job
3 processes a selected one of the at least one ACK rounds.

1 6. A networked computer system comprising:
2 a cluster of computer systems that each includes:
3 a network interface that couples each computer system via a network to
4 other computer systems in the cluster;
5 a memory;
6 a cluster engine residing in the memory that provides a mechanism for
7 communicating ordered messages to and from the computer systems in the cluster;
8 a protocol residing in the memory that specifies at least one data message
9 and at least one acknowledge (ACK) round that provides a time benchmark for
10 determining whether or not a message has been received; and
11 a job residing in the memory and processing the protocol, wherein the job
12 functions according to receiver logic that uses the at least one ACK round to
13 determine without using any timer whether the at least one data message has been
14 received.

1 7. The networked computer system of claim 6 wherein the protocol comprises a
2 plurality of phases that are each followed by an acknowledge (ACK) round.

1 8. The networked computer system of claim 7 wherein each phase of the protocol is
2 defined so that no node can both send a data message and receive a data message during
3 any phase of the protocol.

1 9. The networked computer system of claim 6 wherein the receiver logic comprises
2 post-ACK logic that determines whether the at least one data message has been received
3 after the job processes a selected one of the at least one ACK rounds.

1 10. The networked computer system of claim 6 wherein the receiver logic comprises
2 pre-ACK logic that determines whether the at least one data message has been received
3 before the job processes a selected one of the at least one ACK rounds.

1 11. A computer-implemented method for processing a protocol using a plurality of
2 jobs that form a group in a clustered computing environment, the method comprising the
3 steps of:

4 providing a cluster engine for each member of the group that communicates with
5 the other cluster engines in the group;

6 defining the protocol so that non-receipt of a message by the cluster engine can be
7 determined without using any timer, the protocol including at least one data message and
8 at least one acknowledge (ACK) round that provides a time benchmark for determining
9 whether or not a message has been received; and

10 a job processing the protocol and functioning according to receiver logic that uses
11 the at least one ACK round to determine without using any timer whether the at least one
12 data message has been received.

1 12. The method of claim 11 wherein the step of defining the protocol further
2 comprises the step of defining a plurality of phases that are each followed by an ACK
3 round.

1 13. The method of claim 12 wherein the step of defining the protocol further
2 comprises the step of defining each phase of the protocol so that no node can both send a
3 data message and receive a data message during any phase of the protocol.

1 14. The method of claim 11 further comprising the step of the job using post-ACK
2 logic that determines whether the at least one data message has been received after the job
3 processes a selected one of the at least one ACK rounds.

1 15. The method of claim 11 further comprising the step of the job using pre-ACK
2 logic that determines whether the at least one data message has been received before the
3 job processes a selected one of the at least one ACK rounds.

1 16. A program product comprising:
2 (A) a protocol that specifies at least one data message and at least one
3 acknowledge (ACK) round that provides a time benchmark for determining whether or
4 not a message has been received without using any timer;

5 (B) a job that processes the protocol, wherein the job functions according to
6 receiver logic that uses the at least one ACK round to determine without using any timer
7 whether the at least one data message has been received; and

8 (C) computer-readable signal bearing media bearing the protocol and the job.

1 17. The program product of claim 16 wherein the signal bearing media comprises
2 recordable media.

1 18. The program product of claim 16 wherein the signal bearing media comprises
2 transmission media.

1 19. The program product of claim 16 wherein the protocol comprises a plurality of
2 phases that are each followed by an acknowledge (ACK) round.

1 20. The program product of claim 19 wherein each phase of the protocol is defined so
2 that no node can both send a data message and receive a data message during any phase
3 of the protocol.

1 21. The program product of claim 16 wherein the receiver logic comprises post-ACK
2 logic that determines whether the at least one data message has been received after the job
3 processes a selected one of the at least one ACK rounds.

1 22. The program product of claim 16 wherein the receiver logic comprises pre-ACK
2 logic that determines whether the at least one data message has been received before the
3 job processes a selected one of the at least one ACK rounds.
